arXiv:1601.02461v1 [stat.ME] 11 Jan 2016

Modeling Multivariate Mixed-Response Functional Data

Beth A. Tidemann-Miller[1], Brian J. Reich[2], and Ana-Maria Staicu[2]

January 12, 2016

Abstract

We propose a Bayesian modeling framework for jointly analyzing multiple functional responses of different types (e.g. binary and continuous data). Our approach is based on a multivariate latent Gaussian process and models the dependence among the functional responses through the dependence of the latent process. Our framework easily accommodates additional covariates. We offer a way to estimate the multivariate latent covariance, allowing for implementation of multivariate functional principal components analysis (FPCA) to specify basis expansions and simplify computation. We demonstrate our method through both simulation studies and an application to real data from a periodontal study.

Keywords: Bayesian analysis; Binary data; Markov chain Monte Carlo; Multivariate Functional Principal Components Analysis.

1 Introduction

Until recently, the primary focus of methods employing functional principal components analysis (FPCA) has been on real-valued functional responses. Methods that can model non-Gaussian functional responses, such as repeatedly observed binary or count data, are only recently appearing for univariate functional responses (for example: Hall et al. (2008), van der Linde (2009), Serban et al. (2013)). Additionally, methods that extend functional modeling from the univariate case (i.e. one response curve) to the multivariate case (i.e. a vector of multiple response curves) are currently undergoing development (for example: Zhou et al. (2008), Berrendero et al. (2011), Jacques and Preda (2014)). These multivariate functional methods are limited in that all curves comprising the multivariate response vector must be real-valued.

Here we propose a Bayesian multivariate functional model that utilizes a multivariate latent Gaussian process and can handle responses of different types, e.g. binary and continuous data. Our method easily incorporates

[1]Biogen Inc., Cambridge, Massachusetts
[2]Department of Statistics, North Carolina State University, Raleigh, North Carolina

covariates, a feature previously unavailable for modeling non-Gaussian functional responses. As an extension of the methods of Hall et al. (2008), we propose a way to estimate the multivariate latent covariance, in particular, the cross-covariance of latent functions corresponding to different responses. By using a reliable estimate of the multivariate latent covariance, our proposed method can implement multivariate FPCA to specify basis expansions and simplify computation.

Several approaches to modeling non-Gaussian univariate functional responses have appeared in the literature. For binary or count data observed repeatedly, Hall et al. (2008) proposed a non-parametric functional approach in which the observed responses are directly related to a latent Gaussian functional process through a link function. In order to implement FPCA, they used a Taylor series approximation to derive estimators of the latent process mean function and covariance operator and used bootstrapping methods for further inference. A similar approach by Serban et al. (2013) used logistic functional regression to model multilevel cross-dependent binary-valued functional data. In the case of non-rare events, their approach is an extension of the linear approximation methods of Hall et al. (2008) to multilevel data. For rare events, they introduced an approach centered around an exponential approximation.

In contrast to the aforementioned frequentist methods, van der Linde (2009) offered a Bayesian approach to FPCA for repeatedly observed binary or count data. They extended the variational algorithm for Gaussian responses given in van der Linde (2008), and focused on canonical links for one-parameter exponential families. The methods of Hall et al. (2008), Serban et al. (2013) and van der Linde (2009) offer ways to model univariate functional responses, whereas the approach we propose in this paper jointly models multivariate functional responses of mixed type.

To date, the literature concerning multivariate FPCA has been sparse. Ramsay and Silverman (2005) gave a brief example that uses FPCA for a bivariate functional response of hip and knee angle measurements for gait data. After assigning the two functional responses to a fine grid of points, they concatenated the two response functions and proceeded with PCA in the traditional multivariate framework. Berrendero et al. (2011) proposed multivariate FPCA in which the principal components are smooth functions, a result of performing FPCA at each observed location in a domain on which curves have been smoothed. In contrast to the approach of Ramsay and Silverman (2005), Jacques and Preda (2014) presented a method that al-

lowed for non-orthonormal bases which made it possible for each curve in the multivariate response vector to have its own basis expansion. Their approach neatly addresses how to handle responses with differing magnitudes of variation within the curves.

To our knowledge, our method that models multivariate mixed-type responses is the first of its kind within the functional data analysis literature. In the spatial literature, Reich and Bandyopadhyay (2010) developed a spatial latent factor model for multivariate mixed-response data with informative missingness. Our approach shares several similarities to that of Reich and Bandyopadhyay (2010), however our approach is able to examine complex correlation structures that their stationary spatial method is not equipped to handle.

2 Model

2.1 General Framework

We present the following methodology to jointly model P functional responses. Denote $Y_{pi}(t)$ as the observed functional response of type $p = 1, \ldots, \mathrm{P}$ for subject $i = 1, \ldots, \mathrm{N}$ at location $t \in \mathcal{T}$. The responses $Y_{pi}(t)$ are observed only at a finite set of L_{pi} locations $t_{pi1}, t_{pi2}, \ldots, t_{pi\mathrm{L}_{pi}}$, which may be different for subject and response type. To combine responses with different supports, e.g., binary and continuous, let $Y_{pi}(t) = h_p\{W_{pi}(t)\}$ for link function $h_p(\cdot)$ and latent response $W_{pi}(t)$. Motivated by the periodontal application in Section 5, we restrict our attention to Gaussian and binary responses. If response p is Gaussian then we use the identity link $h_p(\eta) = \eta$; if response p is binary, then we use the indicator link $h_p(\eta) = \mathrm{I}(\eta > 0)$.

Dependence between responses is modeled via the latent Gaussian processes

$$W_{pi}(t) = Z_{pi}(t) + \epsilon_{pi}(t) \tag{1}$$

where $\epsilon_{pi}(t) \overset{iid}{\sim} \mathrm{N}(0, \tau_p^2)$ is random noise and $Z_{pi}(t)$ is a random process. For identification purposes, we fix $\tau_p = 1$ for binary responses. Furthermore, let $Z_{pi}(t) = \mu_{pi}(t) + f_{pi}(t)$, the sum of a fixed mean function $\mu_{pi}(t)$ and a smooth subject-specific process $f_{pi}(t)$, assumed to be uncorrelated with $\epsilon_{pi}(t)$.

The mean can be modeled as $\mu_{pi}(t) = \sum_{j=1}^{\mathrm{m}_p} x_{pij}(t)\beta_{1pj} + s_p(t)$ so that it can incorporate m_p covariates $x_{pij}(t)$ with fixed coefficients β_{1pj} and a

population-level smooth function $s_p(t)$. It is possible for a subject-specific covariate to depend on the functional location t, for example the indicator of jaw in the periodontal data of Section 5, and it is also possible for the same covariates to affect all responses. The smooth function $s_p(t)$ is assumed to be square integrable on $\mathcal{L}^2[0,1]$. We use a predetermined basis expansion to approximate $s_p(t)$. Let $\{\mathrm{B}_{pj}(t) : 1 \leq j \leq \mathrm{n}_p\}$ be a basis expansion in $\mathcal{L}^2[0,1]$ of dimension n_p. We approximate the smooth part by $s_p(t) = \sum_{j=1}^{\mathrm{n}_p} \mathrm{B}_{pj}(t)\beta_{2pj}$ where the type of basis expansions are allowed to differ across response p. To simplify notation, we write $\mu_{pi}(t) = \boldsymbol{u}_{pi}^{\mathrm{T}}(t)\boldsymbol{\beta}_p$ where $\boldsymbol{u}_{pi}(t) = [x_{pi1}(t), \ldots, x_{pi\mathrm{m}_p}(t), \mathrm{B}_{p1}(t), \ldots, \mathrm{B}_{p\mathrm{n}_p}(t)]^{\mathrm{T}}$ is a vector of length $\mathrm{J}_p = \mathrm{m}_p + \mathrm{n}_p$ that combines the covariates and basis functions and has corresponding coefficient vector $\boldsymbol{\beta}_p = [\beta_{1p1}, \ldots, \beta_{1p\mathrm{m}_p}, \beta_{2p1}, \ldots, \beta_{2p\mathrm{n}_p}]^{\mathrm{T}}$.

Let $\boldsymbol{f}_i(t) = [f_{1i}(t), \ldots, f_{\mathrm{P}i}(t)]^{\mathrm{T}}$ be the vector of random subject-specific deviation functions and assume $\boldsymbol{f}_i(t)$ are i.i.d. mean-zero Gaussian processes where $\mathrm{Cov}\{\boldsymbol{f}_i(t), \boldsymbol{f}_i(t')\} = \boldsymbol{K}(t,t')$ and $K_{pp'}(t,t') = \mathrm{Cov}\{f_{pi}(t), f_{p'i}(t')\}$ form the elements of $\boldsymbol{K}(t,t')$. The covariance operator $K_{pp'}(t,t')$ captures both auto-dependence ($p = p'$) and cross-dependence ($p \neq p'$) between two different latent responses. We assume that $f_{pi}(t)$ is a smooth process in $\mathcal{L}^2[0,1]$ and present two ways of specifying basis expansions for $f_{pi}(t)$: Section 2.2 details how to use predetermined bases and Section 2.3 gives a data-driven approach that uses multivariate FPCA.

We can write the multivariate model succinctly in matrix form. Let $\boldsymbol{\beta} = [\boldsymbol{\beta}_1^{\mathrm{T}}, \ldots, \boldsymbol{\beta}_{\mathrm{P}}^{\mathrm{T}}]$ be the fixed effect vector of length $\mathrm{J} = \sum_{p=1}^{\mathrm{P}} \mathrm{J}_p$ with corresponding $\mathrm{P} \times \mathrm{J}$ matrix $\boldsymbol{U}_i(t)$ comprised of appropriate evaluations of $\boldsymbol{u}_{pi}(t)$. Let $\boldsymbol{\epsilon}_i(t) \overset{iid}{\sim} N(0, \boldsymbol{D})$ where $\boldsymbol{D}$ is diagonal with elements $\tau_1^2, \ldots, \tau_{\mathrm{P}}^2$. Then (1) becomes

$$\boldsymbol{W}_i(t) = \boldsymbol{U}_i(t)\boldsymbol{\beta} + \boldsymbol{f}_i(t) + \boldsymbol{\epsilon}_i(t). \tag{2}$$

2.2 Predetermined bases

The first way in which we specify basis expansions for $f_{pi}(t)$ is by choosing predetermined bases such as B-spline, Fourier, or polynomial bases. Let

$$f_{pi}(t) = \sum_{k=1}^{\mathrm{M}_p} \psi_{pk}(t)\alpha_{pik} \tag{3}$$

where $\{\psi_{pk}(t) : 1 \leq k \leq \mathrm{M}_p\}$ is a basis expansion in $\mathcal{L}^2[0,1]$ of dimension M_p and $\boldsymbol{\alpha}_{pi} = [\alpha_{pi1}, \ldots, \alpha_{pi\mathrm{M}_p}]^{\mathrm{T}}$ are random coefficients with $\mathrm{E}(\alpha_{pik}) = 0$ and

$\text{Cov}(\alpha_{pik}, \alpha_{p'i\ell}) = \xi_{k\ell pp'}$. The multivariate covariance function induced by (3) is

$$K_{pp'}(t,t') = \text{Cov}\{f_{pi}(t), f_{p'i}(t')\} = \sum_{k=1}^{\text{M}_p} \sum_{\ell=1}^{\text{M}_{p'}} \psi_{pk}(t)\psi_{p'\ell}(t')\xi_{k\ell pp'}, \qquad (4)$$

which is a function of both the basis functions and covariance $\mathbf{\Sigma} = \{\xi_{k\ell pp'}\}$. Using predetermined basis expansions is extremely flexible; in the Appendix, we discuss how the covariance model can approximate the covariance matrix of any arbitrary finite-dimensional distribution. The choice of M_p is important in that one needs to select a number of basis functions that is sufficient to approximate the covariance well but is not unnecessarily large. We suggest choosing M_p based on a grid search, using criteria such as deviance information criteria (DIC) for comparison.

2.3 Data-driven bases

As an alternative to using predetermined bases, we introduce a novel approach in which we use estimated basis functions that are obtained through FPCA of the multivariate latent covariance. We propose FPCA for multivariate mixed-responses, inspired by Hall et al. (2008) who introduced FPCA for binary-valued functional responses. We too require that the probability of observing a binary event is sufficiently far from zero or one. For simplicity of presentation, we ignore the covariates and discuss how to account for them later in this section.

Recall from (1) that we model the p^{th} response as $Y_{pi}(t) = h_p\{W_{pi}(t)\}$ through the latent Gaussian process $W_{pi}(t) = Z_{pi}(t) + \epsilon_{pi}(t)$ and link function $h_p(\eta)$. Linking the latent response directly to the observed response is equivalent to assuming there is a corresponding monotone link function $g_p(\cdot)$ such that $\text{E}\{Y_{ip}(t)|Z_{pi}(t)\} = g_p\{Z_{pi}(t)\}$; we focus on g_p here. Following Hall et al. (2008), assume that $g_p(\cdot)$ has bounded fourth derivative and that the latent process satisfies $Z_{pi}(t) = \mu_p(t) + \delta X_{pi}(t)$ for fixed mean $\mu_p(t)$, unknown small constant $\delta > 0$, and mean-zero Gaussian random variable $X_{pi}(t)$ that is i.i.d. across subjects i and has both finite variance and finite covariance between $X_{pi}(t)$ and $X_{p'i}(t')$. Our goal is to approximate the latent covariance matrix of $Z_{pi}(t)$ whose covariance operator is $K_{pp'}(t,t') = \text{Cov}\{Z_{pi}(t), Z_{p'i}(t')\}$. Without loss of generality, we restrict our attention to one continuous Gaussian response ($p = 1$) and one binary response ($p = 2$) with link functions

$g_1(\eta) = \eta$ and $g_2(\eta) = \Phi(\eta)$ where $\Phi(\cdot)$ is the standard normal cdf function. For simplicity, we use g to denote g_2 in the following exposition.

The covariance consists of variance components K_{pp} and cross-covariance components $K_{pp'}$. The variance components K_{11} and K_{22} are estimated using the common FPCA for continuous responses Ramsay and Silverman (2002, 2005) as well as binary-valued responses (Hall et al., 2008), respectively. In particular, when the responses are binary valued, the variance K_{22} is estimated using

$$\widetilde{K}_{22}(t,t') = \{\hat{S}_{22}(t,t') - \hat{\eta}_2(t)\hat{\eta}_2(t')\}/[g^{(1)}\{\hat{\mu}_2(t)\}g^{(1)}\{\hat{\mu}_2(t')\}], \tag{5}$$

where $g^{(1)}$ indicates the first derivative of g. The latent mean estimator is $\hat{\mu}_p(t) = g^{-1}\{\hat{\eta}_p(t)\}$ where $\hat{\eta}_p(t)$ estimates $\mathrm{E}[g\{Z_{pi}(t)\}] = \eta_p(t)$ and is found by smoothing the data $\big(t, Y_{pi}(t)\big)$ for $i = 1, \ldots, \mathrm{N}$. $\hat{S}_{22}(t,t')$ is the estimator for $S_{22}(t,t') = \mathrm{E}\{Y_{2i}(t)Y_{2i}(t')\} = \mathrm{E}\big[g\{Z_{2i}(t)\}g\{Z_{2i}(t')\}\big]$ and is obtained through bivariate smoothing of the data $\big((t,t'), Y_{2i}(t)Y_{2i}(t')\big)$ for $i = 1, \ldots, \mathrm{N}$, removing the diagonals before smoothing.

For the cross covariance operator K_{12} we remark that

$$K_{12}(t,t') = \mathrm{Cov}\big\{Y_{1i}(t), Y_{2i}(t')\big\}/g^{(1)}\{\mu_2(t')\}, \tag{6}$$

which is obtained by approximating $\mathrm{Cov}\big\{Y_{1i}(t), Y_{2i}(t')\big\} = \mathrm{Cov}\big[Z_{1i}(t), g\{Z_{2i}(t')\}\big]$ using a Taylor expansion of $g\{Z_{2i}(t')\}$ around $\mu_2(t')$. More details are given in the Appendix. This leads to the estimator of the cross component given by

$$\widetilde{K}_{12}(t,t') = \{\hat{S}_{12}(t,t') - \hat{\eta}_1(t)\hat{\eta}_2(t')\}/g^{(1)}\{\hat{\mu}_2(t')\}. \tag{7}$$

Combining the individually smoothed estimators $\widetilde{K}_{11}(t,t')$, $\widetilde{K}_{22}(t,t')$ and $\widetilde{K}_{12}(t,t') = \widetilde{K}_{21}(t',t)$ forms the smooth 2×2 estimator $\widetilde{\boldsymbol{K}}(t,t')$ of the bivariate latent covariance operator. Note that for smoothing purposes in this paper, we implement a global smoother as opposed to the local least squares smoothing of Hall et al. (2008), though either is appropriate. In the presence of subject-specific covariates, one can find covariate estimates using least squares or logistic regression, depending on the type of response, and then use the residuals to estimate the latent covariance.

The final step for creating basis functions is to implement bivariate FPCA in which we obtain the eigenfunctions $\boldsymbol{\theta}(t) = [\theta_1(t), \ldots, \theta_{\mathrm{P}}(t)]^{\mathrm{T}}$ and the eigenvalues λ of the matrix $\widetilde{\boldsymbol{K}}(t,t')$. Note that the matrix $\widetilde{\boldsymbol{K}}(t,t')$ is not

guaranteed to be positive definite, but we can ensure the truncated spectral decomposition $\widetilde{\boldsymbol{K}}(t,t') = \sum_{k=1}^{\mathrm{M}} \lambda_k \boldsymbol{\theta}_k(t)\{\boldsymbol{\theta}_k(t')\}^{\mathrm{T}}$ is positive definite by restricting the inclusion of only positive eigenvalues and their eigenfunction counterparts. The truncation value M is chosen based on the proportion of variation explained by the eigenvalues as suggested in Di et al. (2009). In particular, specify a cumulative explained variance threshold P_1 and individual explained variance threshold P_2. Define $\mathrm{M} = \min\{k : p_{1k} \geq \mathrm{P}_1, p_{2k} < \mathrm{P}_2\}$ where $p_{1k} = \sum_{i=1}^{k} \lambda_i / \sum_{j=1}^{n} \lambda_j$, $p_{2k} = \lambda_k / \sum_{j=1}^{n} \lambda_j$ and the positive eigenvalues are the first $n \geq k$ eigenvalues. We specify the basis functions to be the eigenfunctions scaled by their associated eigenvalues, $\psi_{pk}(t) = \sqrt{\lambda_k}\theta_{pk}(t)$, and the subject-specific deviation function is approximated by $f_{pi}(t) = \sum_{k=1}^{\mathrm{M}} \psi_{pk}(t)\alpha_{ik}$.

Using this data-driven basis approach, the correlation across responses is largely captured by the basis functions from FPCA. Additionally, since each basis function combines information from all responses, the data-driven approach results in one set of basis functions, eliminating the need to have a set of basis functions for each response. These distinctions offer important advantages over the predetermined basis approach. First, having only one set of basis functions reduces the dimensionality of the random-effect covariance matrix $\boldsymbol{\Sigma}$, making it easier to fit. Second, it allows for further simplification since $\boldsymbol{\Sigma}$ is diagonal. This will offer computational advantages over the predetermined basis method where the burden of capturing the correlation across responses falls entirely on estimating a non-diagonal $\boldsymbol{\Sigma}$ which can potentially have very large dimension.

One important consideration to make when implementing this data-driven basis function approach is to ensure that the variance of the latent process for the continuous component is on a scale similar to that of the latent process for the binary component. We suggest scaling the continuous process by $Y_{1i}(t)/s$ where s is the overall sample standard deviation of the continuous response without regard to t. Since s is a scalar quantity, it is straightforward to scale prior to implementing the latent covariance, FPCA and MCMC estimation algorithms, rescaling only the final results back to the original scale.

2.4 Prior Specification

To complete the Bayesian model, we specify priors for the hyperparameters. The fixed effect parameters $\boldsymbol{\beta}$ are assigned uninformative Gaussian priors.

Let the subject random effect $\boldsymbol{\alpha}_i$ have a Gaussian prior with $\mathrm{Cov}(\boldsymbol{\alpha}_i) = \boldsymbol{\Sigma}$ and assign $\boldsymbol{\Sigma}$ an Inverse Wishart prior. For the error variances of the continuous processes, let τ_p^2 have an uninformative gamma prior; for identifiability τ_p^2 is fixed at 1 for binary processes. In summary,

$$\begin{aligned}
&\boldsymbol{\beta}|\sigma_\beta^2 \sim \mathrm{N_J}(\mathbf{0}, \sigma_\beta^2 \mathbf{I_J}) \\
&\boldsymbol{\alpha}_i|\boldsymbol{\Sigma} \sim \mathrm{N_M}(\mathbf{0}, \boldsymbol{\Sigma}) \\
&\boldsymbol{\Sigma}|q_1, q_2 \sim \mathrm{InvWishart_M}(\boldsymbol{V} = q_2 \mathbf{I_M}, \nu = q_1) \\
&\tau_p^2|l, h \sim \mathrm{InvGamma}(l, h)
\end{aligned} \tag{8}$$

for hyperparameters σ_b^2, q_1, q_2, ℓ, and h, selected to result in weak priors.

3 Computational Details

To facilitate MCMC sampling, we treat the continuous latent processes $W_{pi}(t)$ for binary response as unknown parameters to be updated as part of the sampling as in Albert and Chib (1993). Using this auxiliary variable approach, all parameters have conditional conjugacy due to the prior specifications given in Section 2.4, allowing us to implement Gibbs sampling. The Gibbs sampling algorithm uses full-conditional posteriors derived in the Appendix and which use notation that we now describe.

Denote the observation locations as $t_{pi\ell}$, $\ell = 1, \ldots, \mathrm{L}_{pi}$, for each subject i and response p, giving a total of $\mathrm{L}_i = \sum_{p=1}^{\mathrm{P}} \mathrm{L}_{pi}$ locations. Let $n = \sum_{i=1}^{\mathrm{N}} \mathrm{L}_i$ be the total number of locations observed across all subjects. Let $\boldsymbol{W}_{pi}$ be the vector of length L_{pi} formed by evaluating $W_{pi}(t)$ at every $t_{pi\ell}$. Furthermore, combine $\boldsymbol{W}_{pi}$ for all responses to form one vector $\boldsymbol{W}_i$ of length L_i; $\boldsymbol{U}_i$ and $\boldsymbol{\Psi}_i$ are defined analogously. Then $\boldsymbol{W}_i$ has mean $\mathrm{E}(\boldsymbol{W}_i|\boldsymbol{\alpha}_i) = \boldsymbol{U}_i\boldsymbol{\beta} + \boldsymbol{\Psi}_i\boldsymbol{\alpha}_i$ and precision matrix $\boldsymbol{P}_i$ is comprised of the appropriate error variance parameter τ_p^{-2}.

MCMC begins by setting initial values for all parameters and then sequentially sampling each parameter conditioned on all the others (denoted by "$|\cdot$"). Sampling is performed (using the latest sample to update each parameter) according to the full conditionals in the following manner:

1. Select initial values for $\boldsymbol{\beta}$, $\boldsymbol{\alpha}_i$, $\boldsymbol{\Sigma}$, $W_{pi}(t)$ for binary responses, and τ_p^2 for continuous responses;

2. For each $i = 1, \ldots, \mathrm{N}$ and $\ell = 1, \ldots, \mathrm{L}_{pi}$, update the latent response corresponding to the observed binary response by drawing from $W_{pi}(t_{i\ell})|\cdot \sim \mathrm{N}(\boldsymbol{u}_{pi}^{\mathrm{T}}(t_{i\ell})\boldsymbol{\beta} + \boldsymbol{\psi}_p^{\mathrm{T}}(t_{i\ell})\boldsymbol{\alpha}_i, 1)$ restricted to the interval $(0, \infty)$ if $Y_{pi}(t_{i\ell}) = 1$ or $(-\infty, 0)$ if $Y_{pi}(t_{i\ell}) = 0$;

3. Update the population mean parameter by drawing from $\boldsymbol{\beta}|\cdot \sim \mathrm{N}(\boldsymbol{\mu}_{\boldsymbol{\beta}}, \boldsymbol{V}_{\boldsymbol{\beta}})$ for $\boldsymbol{V}_{\boldsymbol{\beta}} = \left[\left(\sum_{i=1}^{\mathrm{N}} \boldsymbol{U}_i^{\mathrm{T}} \boldsymbol{P}_i \boldsymbol{U}_i\right) + \sigma_{\boldsymbol{\beta}}^{-2} \mathbf{I}_{\mathrm{J}}\right]^{-1}$ and $\boldsymbol{\mu}_{\boldsymbol{\beta}} = \boldsymbol{V}_{\boldsymbol{\beta}} \left[\sum_{i=1}^{\mathrm{N}} \boldsymbol{U}_i^{\mathrm{T}} \boldsymbol{P}_i (\boldsymbol{W}_i - \boldsymbol{\Psi}_i \boldsymbol{\alpha}_i)\right]$;

4. For each $i = 1, \ldots, \mathrm{N}$, update the random effect by sampling from $\boldsymbol{\alpha}_i|\cdot \sim \mathrm{N}(\boldsymbol{\mu}_{\boldsymbol{\alpha}}, \boldsymbol{V}_{\boldsymbol{\alpha}})$ for $\boldsymbol{V}_{\boldsymbol{\alpha}} = \left(\boldsymbol{\Psi}_i^{\mathrm{T}} \boldsymbol{P}_i \boldsymbol{\Psi}_i + \boldsymbol{\Sigma}^{-1}\right)^{-1}$ and $\boldsymbol{\mu}_{\boldsymbol{\alpha}} = \boldsymbol{V}_{\boldsymbol{\alpha}} \boldsymbol{\Psi}_i^{\mathrm{T}} \boldsymbol{P}_i (\boldsymbol{W}_i - \boldsymbol{U}_i \boldsymbol{\beta})$;

5. Update the random effect covariance matrix through $\boldsymbol{\Sigma}|\cdot \sim \mathrm{InvWishart}_{\mathrm{M}}[\{\sum_{i=1}^{\mathrm{N}} \boldsymbol{\alpha}_i \boldsymbol{\alpha}_i^{\mathrm{T}} + (1/q_2)\mathbf{I}_{\mathrm{M}}\}^{-1}, \mathrm{N} + q_1]$;

6. Update the error variance for the continuous responses according to $\tau_p^2|\cdot \sim \mathrm{InvGamma}(l_\omega, h_\omega)$ with $l_\omega = n/2 + l$ and $h_\omega = h + 1/2 \sum_{i=1}^{\mathrm{N}} \sum_{\ell=1}^{\mathrm{L}_i} [W_{pi}(t_{i\ell}) - \boldsymbol{u}_{pi}^{\mathrm{T}}(t_{i\ell})\boldsymbol{\beta} + \boldsymbol{\psi}_p^{\mathrm{T}}(t_{i\ell})\boldsymbol{\alpha}_i]^2$.

Steps 2-6 are repeated for the desired number of samples.

4 Simulations

For our simulation study, we compare mean estimation and prediction performance among four estimating models when the generating model has a continuous and binary response, either generated separately (univariate) or jointly (bivariate) with strong cross-correlation, for both small and large sample size. The four estimating models are either univariate models applied separately to each response or bivariate models, and either employ the prespecified basis function method of Section 2.2 or the data-driven approach of Section 2.3.

4.1 Data generation

We consider the case where $Y_{1i}(t)$ is continuous and $Y_{2i}(t)$ is binary. Functions are observed at a dense, balanced design with $\mathrm{L}_{pi} \equiv 30$ equally-spaced locations in $[0, 1]$ for each subject i and response p. We use the model given

in (2) with predetermined bases as in Section 2.2 for data generation. We specify a separable random effect covariance matrix $\boldsymbol{\Sigma} = \boldsymbol{A} \otimes \boldsymbol{C}$, where $\text{Cov}([\alpha_{1ik}, \alpha_{2ik}]^{\text{T}}) = \boldsymbol{A}$ for $\boldsymbol{A}_{11} = \boldsymbol{A}_{22} = 1$ and $\boldsymbol{A}_{12} = \boldsymbol{A}_{21} = \rho_{\boldsymbol{\alpha}}$ so that the parameter $\rho_{\boldsymbol{\alpha}}$ controls the correlation between the latent responses, and $\text{Cov}([\alpha_{pi1}, \ldots, \alpha_{pi\text{M}}]^{\text{T}}) = \boldsymbol{C}$ for $p = 1, 2$ controls the covariance of the random effect basis function coefficients and is the same across responses. The $\boldsymbol{C}$ used for data generation has the AR(1) structure with variance 1 and correlation parameter $\rho = 1/2$.

For the fixed population mean function we assume there are no subject-level covariates so that $\boldsymbol{\mu}(t) = \mathbf{B}(t)\boldsymbol{\beta}$, and we specify a quadratic basis $\{\text{B}_{pj}(t) = t^{(j-1)} : 1 \leq j \leq 3\}$ for each response p with coefficients $\boldsymbol{\beta}_1 = [-0.64, 4, -4]^{\text{T}}$ and $\boldsymbol{\beta}_2 = [0.97, -6, 6]^{\text{T}}$. The intercepts are chosen such that the curves are positive for approximately half of the observed locations t. The basis functions for the subject-specific deviation function $\boldsymbol{f}_i(t) = \boldsymbol{\Psi}(t)\boldsymbol{\alpha}_i$ are given by $\psi_{1k}(t) = \sin\{(2\pi k/\text{M})(t + 2\pi k/\text{M})\}$ and $\psi_{2k}(t) = \cos\{(2\pi k/\text{M})(t + 2\pi k/\text{M})\}$ for $k = 1, \ldots, \text{M} = 7$. The error variance for the continuous process is $\tau_1^2 = 1$. We generate data from four scenarios given in Table 1 by varying the sample size ($\text{N} = 50, 250$) and the cross-correlation ($\rho_{\boldsymbol{\alpha}} = 0, 0.8$). All scenarios use 100 Monte Carlo (MC) replications.

4.2 Models and metrics for comparison

We fit four models to each dataset.

I. Bivariate B-spline (BBSP): the multivariate model in (2) with B-spline bases as in Section 2.2;

II. Univariate B-spline (UBSP): the model from (1) applied separately to each response with B-spline bases as in Section 2.2;

III. Bivariate FPCA (BFPCA): the multivariate model in (2) with data-driven bases as in Section 2.3;

IV. Univariate FPCA (UFPCA): the model in (1) applied separately to each response with data-driven bases as in Section 2.3;

For estimation using the B-spline methods, we choose B-splines of order 4 and the number of B-spline breaks for each replication is fixed at 6 based on preliminary analyses. For the FPCA methods, we specify an unstructured $\boldsymbol{\Sigma}$, and the number of basis functions is chosen to explain at least $\text{P}_1 = 99\%$ of

the cumulative variation. In practice, both the number of basis functions for the B-spline method and the percentage of variation explained for the FPCA method are tuning parameters and one should compare results over a grid parameter values. For the population mean we fit the true polynomial basis $\mathbf{B}(t)$ for estimation. We perform MCMC sampling with 20,000 draws and the first 5,000 are discarded as burn-in. The hyperparameters are specified as $\sigma_b^2 = 100$ and $q_1 = q_2 = l = h = 0.1$.

Methods are compared in terms of their predictive performance and ability to estimate the marginal mean function for each response. Let $\omega_{1i}(t) = \mathrm{E}\{Y_{1i}(t)\} = \boldsymbol{u}_{1i}^{\mathrm{T}}(t)\boldsymbol{\beta}_1$ and $\omega_{2i}(t) = \mathrm{E}\{Y_{2i}(t)\} = \Phi\{\gamma_i(t)\}$, where $\gamma_i(t) = \boldsymbol{u}_{2i}^{\mathrm{T}}(t)\boldsymbol{\beta}_2/\sqrt{v_2(t)}$ is the population effect shrunk toward zero by the square root of the marginal variance $v_2(t) = \mathrm{Var}\{Y_2(t)\} = \boldsymbol{\psi}_2(t)\boldsymbol{\Sigma}_{22}\{\boldsymbol{\psi}_2(t)\}^{\mathrm{T}} + 1$. Let $\widehat{\omega}_{pr}(t)$ and $\hat{\nu}_{pr}(t)$ be the posterior mean and variance, respectively, for MC replication $r = 1, \ldots, 100$. Metrics for comparison of estimated means found in Table 1 for each response are mean integrated squared error: $\mathrm{MISE} = \int_t \mathrm{E}\{\widehat{\omega}_p(t) - \omega_p(t)\}^2 dt$; coverage of 95% pointwise confidence intervals $\widehat{\omega}_{pr}(t) \pm l_{pr}(t)$ averaged over location t and MC replication r with margin of error $l_{pr}(t) = 1.96\sqrt{\hat{\nu}_{pr}(t)}$; and confidence interval length $2l_{pr}(t)$.

For prediction, we generate additional data $Y_{prj}(t_l)$ at equally spaced locations $t_\ell \in [0, 1]$ where $\ell = 1, \ldots, 30$ for subjects $j = 1, \ldots, 20$ per response $p = 1, 2$ for each MC replication $r = 1, \ldots, 100$. To assess the value of jointly modeling the two responses, we leave out all of response 1 for 10 subjects and all of response 2 for the remaining 10 subjects per replication. Models are compared in terms of their predictive performance using mean squared prediction error (MSPE) for each response, defined as $\mathrm{MSPE} = (nm\mathrm{L})^{-1} \sum_{r=1}^{n} \sum_{j=1}^{m} \sum_{\ell=1}^{\mathrm{L}} \{Y_{prj}(t_\ell) - \hat{Y}_{prj}(t_\ell)\}^2$. For binary responses this is known as the Brier score and $\hat{Y}_{prj}(t_\ell)$ is the posterior probability that $Y = 1$.

4.3 Results

Table 1 gives the simulation results. There appears to be little difference in mean function estimation between univariate and bivariate methods for all scenarios. When strong correlation is present (Scenarios 1 & 2), the bivariate methods show marked improvement in prediction for both responses over the univariate methods, a difference that becomes more pronounced with an increase in sample size. Bivariate methods perform well when the generating model is univariate (Scenarios 3 & 4). Though prediction is better when

Table 1: Simulation Results

	Continuous Response				Binary Response			
	MISE	CI length	95 % Cvg	MSPE	MISE	CI length	95 % Cvg	MSPE
Scenario 1: $n = 50$, $\rho_{\boldsymbol{\alpha}} = 0.8$								
BFPCA	3.50	65.3	92.9 ***	319	0.174	12.8	86.9 ***	23.0
BBSP	3.26	61.6 **	90.7 ***	313	0.168	12.9	87.3 ***	22.9
UFPCA	3.20	66.5 *	93.6	350 *	0.182	14.8 *	90.8 ***	24.4 *
UBSP	2.86	64.9	94.0	351 *	0.185	14.5 *	90.9 ***	24.3 *
Scenario 2: $n = 250$, $\rho_{\boldsymbol{\alpha}} = 0.8$								
BFPCA	0.795	31.9	91.4 ***	284	0.039	6.66	90.7 ***	21.0
BBSP	0.798	31.2 **	91.0 ***	285	0.037	6.67	91.3 ***	21.0
UFPCA	0.790	32.8 *	91.9 ***	351 *	0.040	7.27 *	93.2	24.3 *
UBSP	0.794	32.4 *	92.0 ***	350 *	0.043	7.25 *	91.6 ***	24.3 *
Scenario 3: $n = 50$, $\rho_{\boldsymbol{\alpha}} = 0$								
BFPCA	2.96	65.9	94.2	408	0.172	13.4	89.3 ***	26.3
BBSP	3.16	62.6 **	92.8 ***	421 *	0.183	13.3	88.4 ***	26.6 *
UFPCA	2.75	65.8	94.6	372 **	0.166	14.6 *	93.3	24.4 **
UBSP	2.85	63.9 **	93.5	371 **	0.162	14.5 *	92.8 ***	24.2 **
Scenario 4: $n = 250$, $\rho_{\boldsymbol{\alpha}} = 0$								
BFPCA	0.802	32.5	94.6	370	0.044	6.96	91.1 ***	24.8
BBSP	0.791	31.9 **	94.3	374 *	0.042	6.97	90.7 ***	24.9
UFPCA	0.780	32.9 *	94.7	362 **	0.042	7.35 *	93.5	24.3 **
UBSP	0.765	32.5	94.5	361 **	0.040	7.35 *	94.1	24.3 **

Results in hundredths. A '**' ('*') indicates better (worse) compared to BFPCA by Wilcoxson rank sum test, $\alpha = 0.05$. For coverage, a '***' indicates that the coverage is not within the nominal 95% range.

fitting the correct univariate model, the differences between the bivariate and univariate methods become very small with an increase in sample size. All methods show slight under-coverage.

For Scenarios 1 & 2 there is no clear difference between fitting predetermined bases (Section 2.2) or data-driven bases (Section 2.3); however, BFPCA has better prediction compared to BBSP in Scenarios 3 & 4 when there is no cross-correlation. The univariate models have very similar performance to one another in all scenarios.

5 Periodontal Data Application

We demonstrate our methods using data from a periodontal study (Fernandez et al., 2009) conducted by the Center for Oral Health Research at the Medical University of South Carolina. In addition to collecting subject-level covariates for over 200 Gullah African Americans, several measures of patients' periodontal health were observed at six sites for each of 28 teeth. The two responses we consider are (continuous) clinical attachment loss (CAL) and (binary-valued) bleeding on probing (BOP). CAL is the distance that a tooth has detached from the bone, rounded to the nearest mm. We use the average CAL over the six sites on each tooth as the tooth's CAL response. BOP is the binary indicator of whether the gums bleed when pressed with a dental probe at any of the six sites per tooth. A total of N = 197 patients (subjects) are included for analysis after excluding those with more than 50% missingness. Any remaining missingness is assumed to be completely at random; Reich and Bandyopadhyay (2010) and Reich et al. (2013) provide methods for accounting for non-random missingness.

For our analysis, we assign teeth the numbers 1-14 going from left to right in the upper jaw when looking at a patient and 15-28 going from right to left in the lower jaw when looking at a patient; wisdom teeth are excluded. Using this numbering system, teeth 1 & 28 are adjacent going from upper jaw to lower jaw, and it is the same for teeth 14 & 15 on the other side of the mouth. We consider responses at each tooth to be realizations of a functional process with locations $t \in [1, 28]$. In fitting a bivariate functional model to this data, we hope to gain a better understanding of the dynamics between the responses CAL and BOP through close examination of their cross-covariance. Our extremely flexible approach to modeling the covariance will be able to capture any spatial correlation of adjacent teeth, of teeth on different sides of the mouth, and of teeth on different jaws.

The subject-specific covariates that we include in modeling the mean function are the same covariates used by Reich and Bandyopadhyay (2010) and include age (in years), gender (female=1, male=0), body mass index or BMI (in kg/m^2), smoking status (1=smoker, 0=never), and glycosylated hemoglobin or HbA1c (1 = high, 0 = controlled). All covariates have been standardized to be zero-mean with standard deviation of 1. For each tooth, we include an indicator of jaw (0=upper, 1=lower). For the smooth part of the mean, we consider a quadratic function $s_p(t) = \beta_{p0} + \beta_{p1} d(t) + \beta_{p2} d(t)^2$ of tooth distance d from the front of the mouth, where $d(t) = t - 7.5$ for teeth

Table 2: Model comparisons for the periodontal data application

Model	Subject RE	PVE	PCA	Dbar	pD	DIC
1	Y	99	B	8114	1257	9371
2	Y	99	U	8228	1231	9459
3	Y	95	B	8849	1001	9850
4	Y	95	U	8536	1114	9650
5	N	99	B	10101	942	11043
6	N	99	U	10586	832	11418
7	N	95	B	10511	788	11299
8	N	95	U	10722	760	11482

"Subject RE" indicates inclusion of a subject-specific random intercept. "PVE" is the threshold for cumulative percentage of variation explained. "PCA" indicates whether univariate ("U") or bivariate ("B") FPCA was performed.

in the upper jaw and $d(t) = t - 21.5$ for the lower jaw.

We present analysis for 8 models given in Table 2 that all employ the data-driven basis method of Section 2.3. The 8 models differ by: 1) whether FPCA is univariate or bivariate; 2) the choice of threshold $\mathrm{P}_1 = 99\%, 95\%$ for the cumulative percentage of variation explained for FPCA; and 3) whether a random bivariate subject-level intercept $\boldsymbol{\alpha}_{0i} = [\alpha_{01i}, \alpha_{02i}]^{\mathrm{T}}$ is added to model (2). Models using B-splines as in Section 2.2 were also considered but are not presented because the best-performing models required a large number of basis functions.

For the purpose of estimating the latent covariance, we ignore the covariates. When incorporating a bivariate random subject-level intercept, we use residuals $R_{1i}(t) = Y_{1i}(t) - \mathrm{L}_{1i}^{-1} \sum_{i=1}^{\mathrm{L}_{1i}} Y_{1i}(t)$ of the continuous response CAL to estimate the latent covariance for FPCA; this is not done for the binary responses as the residuals would no longer be binary. For models that include $\boldsymbol{\alpha}_{0i}$, we estimate the covariance term $\mathrm{Cov}(\alpha_{01i}, \alpha_{02i})$ in addition to the variance terms $\mathrm{Var}(\alpha_{0pi})$. We specify a diagonal covariance matrix $\boldsymbol{\Sigma}$ for the remaining random effect parameters.

Table 2 shows that Models 1-4, which include a subject random effect, outperform (based on DIC) Models 5-8 which omit the subject random effect. For this data, specifying the larger percentage of variation explained for FPCA, and hence including more basis functions, leads to better model performance. In comparing the two leading models 1 & 2, implementing FPCA

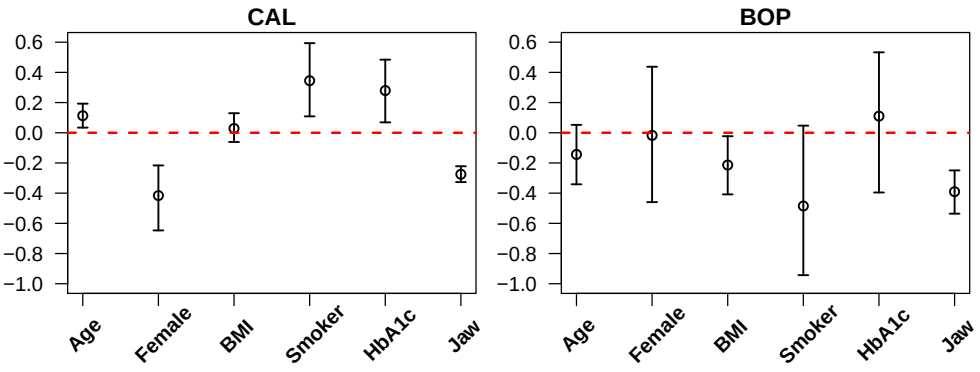

Figure 1: Posterior medians and 95% posterior intervals of the subject-specific covariate coefficients by response.

on the full bivariate covariance matrix as in Model 1, taking into account the cross-dependence between the two responses CAL and BOP, leads to superior performance. Figure 1 shows the subject-level coefficient estimates and 95% posterior intervals for Model 1. Models 2-4 had similar coefficient estimates. For CAL, only the coefficient interval for BMI includes zero. The other coefficient estimates show an increased level of CAL for older patients, males, smokers, patients with high HbA1c counts, and for teeth on the upper jaw. For BOP, the posterior confidence intervals are larger than those for CAL. For intervals that exclude zero, there is an increase of BOP for the upper jaw, yet a slightly lower incidence of BOP for higher BMI.

Figure 2 shows the fitted values (from Model 1) for two individuals in the periodontal data set. The left panels show the posterior means and 95% posterior intervals of the subject-specific mean function $\mu_{1i}(t)$ for the continuous response CAL. Most of the observed CAL values fall within the 95% interval for both subjects, indicating a reasonable model fit. The right panels show the posterior mean and 95% posterior intervals of the conditional probability of the event, $P(Y_{2i}(t) = 1|\boldsymbol{\alpha}_{2i})$. Teeth with observed BOP $(= 1)$ are indicated by the squares on the bottom of the plot. The higher predicted probabilities tend to correspond to the incidence of BOP, again indicating a reasonable model fit.

The posterior summaries of the auto- and cross-correlations of the subject-specific process $\boldsymbol{f}_i(t)$ from (2) are given in Figure 3; note that the

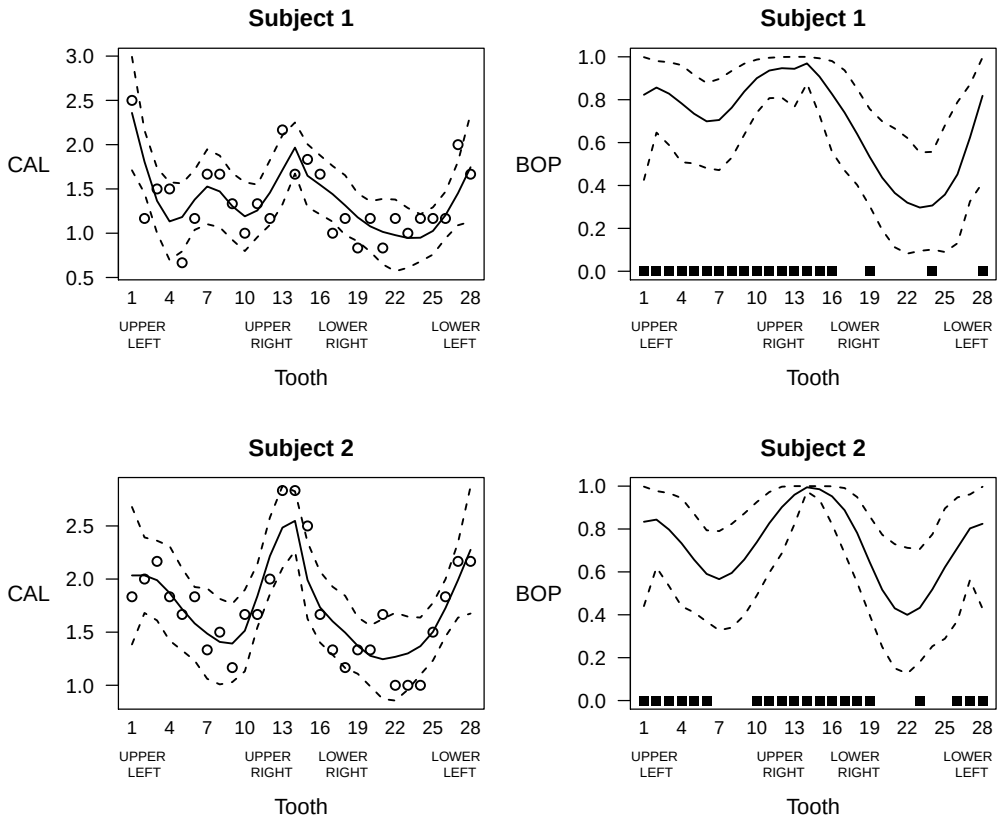

Figure 2: Fitted values for two individuals from the periodontal study (using Model 1). *Left panels*: Observed values of CAL are shown as dots. The solid black line indicates the posterior mean of $\mu_{1i}(t)$, the subject-specific mean function, and point-wise 95% posterior intervals are given by the dotted lines. *Right panels*: The squares along the x-axis indicate the teeth for which BOP is observed. The solid black line gives the posterior mean of the conditional probability of the event, $P(Y_{2i}(t) = 1|\boldsymbol{\alpha}_{2i})$, and dotted lines show point-wise 95% posterior intervals. The label "UPPER LEFT" refers to the left side of the the upper jaw when looking at a patient, and it is analogous for the other labels.

correlation attributed to the subject random intercept is not included in this figure. In this periodontal application, these plots offer important and novel insights into the complex relationships that exist between and within the BOP and CAL responses in different parts of the mouth. The utility of quantifying and visualizing these complex correlation relationships is apparent for many other types of applications.

Examination of the diagonal of the auto-correlation plot for CAL in Figure 3 shows strong positive spatial correlation between adjacent teeth and between teeth separated by only one or two teeth on the same jaw. This plot also shows positive correlation between a tooth in the left and a tooth in the right side of the same jaw, and the relationship is particularly strong for teeth in the lower jaw. The correlation for CAL between teeth in opposite sides of the mouth and on different jaws is also positive, yet not as strong as for teeth on the same jaw; this correlation is very similar in magnitude as the correlation for teeth on the left or right side of the mouth but on different jaws. Additionally, there are mild to strong negative correlations between teeth in the center (front) of the mouth and teeth in the back of the mouth, regardless of the jaws on which the teeth are located. This is also seen in the plot of the posterior probability that the auto-correlation is positive.

In the auto-correlation plot for BOP, again we see strong positive spatial correlation between adjacent teeth and between teeth that are close to one another on the same jaw. Additionally, the plots of the auto-correlation and of the probability of being positive show that the correlation is mostly positive with only a few areas of negative correlation. The correlation is negative between a tooth in the center and a tooth on the right side of the upper jaw, as well as between a tooth in the left and a tooth in the center of the lower jaw. There is also a strong negative correlation for teeth in the lower right and upper right, as well as for teeth in the lower left and lower right.

The cross-correlation between BOP and CAL ranges from moderately positive to moderately negative. Unlike the auto correlation plots, the cross correlation is not symmetric, which makes interpretation slightly more complex. For instance, BOP in the lower left is positively correlated with CAL in the center and upper right as indicated by the darkest patch near the top center of the cross correlation figure. Alternatively, CAL in the lower left shows slightly negative to no correlation with BOP in the center and upper right of the mouth. Another demonstration of this non-symmetric property occurs for the negative correlation of BOP in the lower left with CAL in

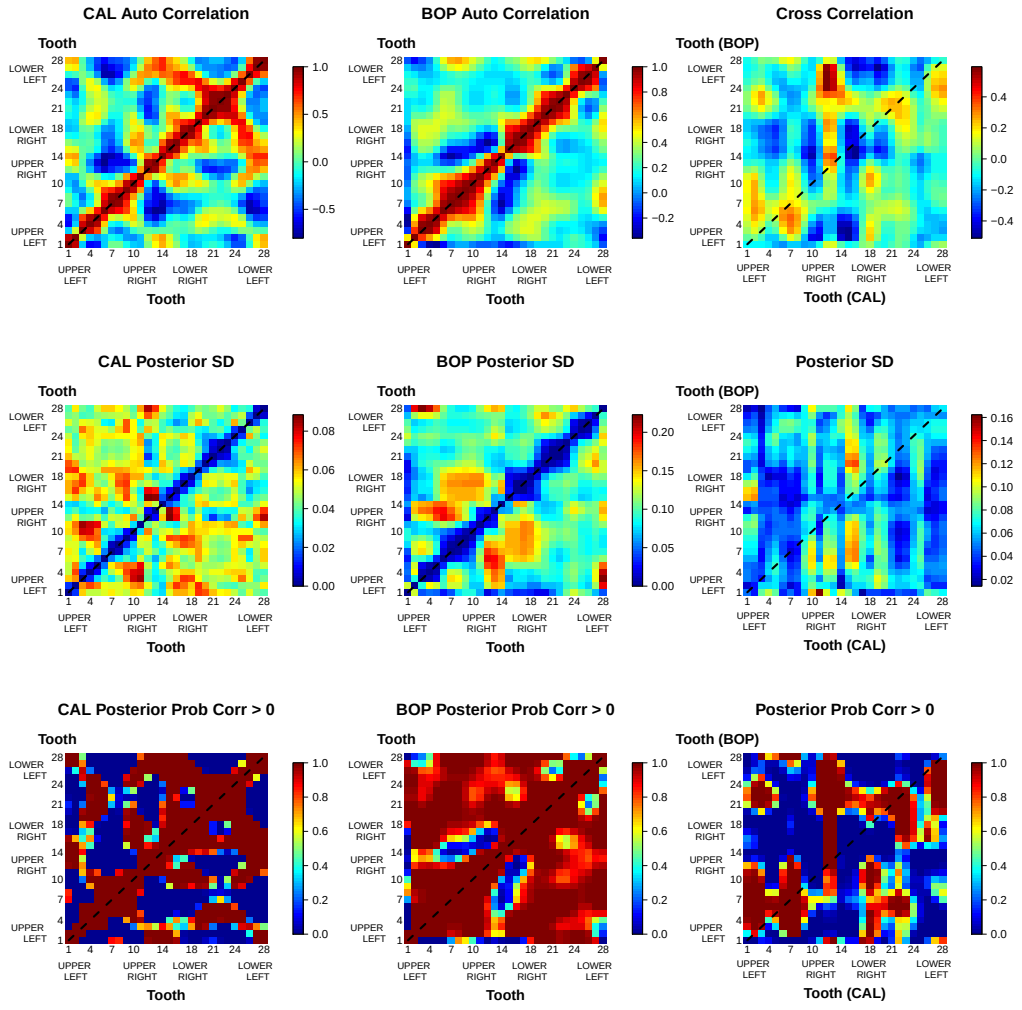

Figure 3: Posterior summaries of the within-response and between-response correlation structures for any two teeth when fit with Model 1 (excluding correlation from the subject random intercepts). The label "UPPER LEFT" refers to the left side of the the upper jaw when looking at a patient, and it is analogous for the other labels.

the lower right, though BOP in the lower right shows slightly positive to no correlation with CAL in the lower left.

6 Discussion

We introduce a methodology to jointly model multivariate functional responses of mixed type (e.g. continuous and binary data) and also propose an extension of FPCA for mixed-responses. Our method can account for subject-specific covariates that can be either linear or time-dependent (such as the jaw indicator used in the analysis of the periodontal study in Section 5). The proposed method is flexible enough for functions to be observed at varying locations for different subjects and different responses. For exposition we focus on modeling a bivariate response vector where one functional response is continuous and the other is binary, though joint modeling of more than two responses is a straightforward extension. Furthermore, the method easily models repeatedly observed categorical responses. This is achieved in a manner similar to thresholding the latent process at zero for binary data, but instead one must impose multiple thresholds on the latent process. Modeling other types of data, such as repeatedly observed count data, is not as straightforward as it would likely require using copulas

By estimating the multivariate covariance of the latent process, our methodology can offer novel insights into the cross-dependence of different responses, which is of interest in a wide variety of applications. Quantifying and exploring this dependence is an important contribution of our method and is a primary goal of our analysis of the periodontal data presented in Section 5. Reich and Bandyopadhyay (2010) and Reich et al. (2013) offer ways to incorporate informative missingness and apply their methods to the same periodontal data. We do not address the informative missingness for our analysis because it is not central to our goals, and leave it for future work.

Acknowledgments

The authors thank Dipankar Bandypadhyay of the University of Minnesota and Drs. S. London, J. Fernandes, C. Salinas, W. Zhao, Ms. L. Summerlin and Ms. P. Hudson. W of the Center for Oral Health Research (COHR) at

the Medical University of South Carolina for providing the data and context for this work. The work of Beth A. Tidemann-Miller was supported by NIH grant R01 CA085848.

References

Albert, J. H. and Chib, S. (1993) Bayesian analysis of binary and polychotomous response data. *Journal of the American Statistical Association*, **88**, pp. 669–679.

Berrendero, J. R., Justel, A. and Svarc, M. (2011) Principal components for multivariate functional data. *Computational Statistics & Data Analysis*, **55**, 2619–2634.

Di, C.-Z., Crainiceanu, C. M., Caffo, B. S. and Punjabi, N. M. (2009) Multilevel functional principal component analysis. *Annals of Applied Statistics*, **3**, 458–488.

Fernandez, J. K., Wiegand, R. E., Salinas, C. F., Grossi, S. G., Sanders, J. J., Lopes-Virella, M. F. and Slate, E. H. (2009) Periodontal disease status in Gullah African Americans with Type 2 diabetes living in South Carolina. *Journal of Periodontology*, **80**, 1062–8.

Hall, P., Müller, H.-G. and Yao, F. (2008) Modelling sparse generalized longitudinal observations with latent Gaussian processes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **70**, 703–723.

Jacques, J. and Preda, C. (2014) Model-based clustering for multivariate functional data. *Computational Statistics & Data Analysis*, **71**, 92 – 106.

van der Linde, A. (2008) Variational bayesian functional {PCA}. *Computational Statistics & Data Analysis*, **53**, 517 – 533.

— (2009) A bayesian latent variable approach to functional principal components analysis with binary and count data. *AStA Advances in Statistical Analysis*, **93**, 307–333.

Ramsay, J. and Silverman, B. (2002) *Applied Functional Data Analysis: Methods and Case Studies*. New York: Springer.

— (2005) *Functional Data Analysis*. New York: Springer, 2 edn.

Reich, B. J. and Bandyopadhyay, D. (2010) A latent factor model for spatial data with informative missingness. *The Annals of Applied Statistics*, **4**, 439–459.

Reich, B. J., Bandyopadhyay, D. and Bondell, H. D. (2013) A nonparametric spatial model for periodontal data with nonrandom missingness. *Journal of the American Statistical Association*, **108**, 820–831.

Serban, N., Staicu, A.-M. and Carroll, R. J. (2013) Multilevel cross-dependent binary longitudinal data. *Biometrics*, **69**, 903–913.

Zhou, L., Huang, J. Z. and Carroll, R. J. (2008) Joint modelling of paired sparse functional data using principal components. *Biometrika*, **95**, 601–619.

Appendix

Appendix A

In this section, we show that we can approximate any smooth covariance using the predetermined basis method. For simplicity, assume that the functional responses are observed at the same locations $t_{p\ell} \equiv t_\ell$ for $\ell = 1, \ldots, \mathrm{L}$ for each response p. We specify this model for an arbitrary subject, and thus drop the subscript i. Let $\boldsymbol{\psi}_{pk}$ be the vector of length L formed by evaluating at every t_ℓ the basis functions $\psi_{pk}(t)$, $k = 1, \ldots, \mathrm{M}_p$, and define the vector $\boldsymbol{f}_p$ analogously. Then we form the $\mathrm{L} \times \mathrm{M}_p$ matrix $\boldsymbol{\Psi}_p = [\boldsymbol{\psi}_{pk}, \ldots, \boldsymbol{\psi}_{p\mathrm{M}_p}]$ and the coefficient vector $\boldsymbol{\alpha}_{pi} = [\alpha_{p1}, \ldots, \alpha_{p\mathrm{M}_p}]^{\mathrm{T}}$ so that we can write $\boldsymbol{f}_p = \boldsymbol{\Psi}_p \boldsymbol{\alpha}_p$. We combine $\boldsymbol{f}_p$ for all responses to form one vector $\boldsymbol{f}$ of length $\mathrm{n} = \mathrm{PL}$, and define the coefficient vector $\boldsymbol{\alpha}^{\mathrm{T}} = [\boldsymbol{\alpha}_1^{\mathrm{T}}, \ldots, \boldsymbol{\alpha}_{\mathrm{P}}^{\mathrm{T}}]$ of length $\mathrm{m} = \sum_{p=1}^{\mathrm{P}} \mathrm{M}_p$ and corresponding block-diagonal matrix $\boldsymbol{\Psi}$ of dimension $\mathrm{n} \times \mathrm{m}$ with blocks $\boldsymbol{\Psi}_p$. The resulting vector $\boldsymbol{f} = \boldsymbol{\Psi}\boldsymbol{\alpha}$ has length n, and we assume $\boldsymbol{\alpha} \overset{iid}{\sim} N(0, \boldsymbol{\Sigma})$ where $\boldsymbol{\Sigma}$ is a covariance matrix of dimension $\mathrm{m} \times \mathrm{m}$ with elements $\mathrm{Cov}(\alpha_{pk}, \alpha_{p'\ell}) = \xi_{k\ell pp'}$.

To illustrate the flexibility of the model, assume $\boldsymbol{\Omega}_0$ is the true $\mathrm{n} \times \mathrm{n}$ covariance matrix of $\boldsymbol{f}$ evaluated at locations t_l. $\boldsymbol{\Omega}_0$ is now approximated by the variance-covariance matrix $\boldsymbol{\Omega} = \mathrm{Cov}(\boldsymbol{\Psi}\boldsymbol{\alpha}) = \boldsymbol{\Psi}\boldsymbol{\Sigma}\boldsymbol{\Psi}^{\mathrm{T}}$. Since the basis

comprising $\boldsymbol{\Psi}$ is pre-specified, the quality of the approximation $\boldsymbol{\Omega} \approx \boldsymbol{\Omega}_0$ is reliant on $\boldsymbol{\Sigma}$. By fitting a large number of basis functions, i.e. setting m = n, it is possible to fit any smooth covariance function. When m = n then $\boldsymbol{\Psi}_i$ is a square matrix. Assume $\boldsymbol{\Psi}_i$ is full rank and thus $\boldsymbol{\Psi}_i^{\mathrm{T}}\boldsymbol{\Psi}_i$ is invertible. Pre- and post- multiplication gives $\boldsymbol{\Psi}_i^{\mathrm{T}}\boldsymbol{\Omega}\boldsymbol{\Psi}_i = \boldsymbol{\Psi}_i^{\mathrm{T}}\boldsymbol{\Psi}_i\boldsymbol{\Sigma}\boldsymbol{\Psi}_i^{\mathrm{T}}\boldsymbol{\Psi}_i$. Since $\boldsymbol{\Theta} = \{\boldsymbol{\Psi}_i^{\mathrm{T}}\boldsymbol{\Psi}_i\}^{-1}$ exists we can recover $\boldsymbol{\Sigma} = \boldsymbol{\Theta}\boldsymbol{\Psi}_i^{\mathrm{T}}\boldsymbol{\Omega}\boldsymbol{\Psi}_i\boldsymbol{\Theta}$. Though this approach is quite flexible, it is hard to estimate $\boldsymbol{\Sigma}$ if it is high-dimension; therefore it is unlikely to perform well if the processes cannot be represented by a small number of basis functions.

Appendix B

Here we describe in more detail the derivation of the latent cross covariance estimator. As our approach is inspired by Hall et al. (2008), we start with a brief summary of the method they proposed for finding the auto-covariance of the latent process corresponding to the binary response, that is, response $p = 2$. First, estimate the mean function for $p = 2$, $\hat{\mu}_2(t) = g^{-1}\{\hat{\eta}_2(t)\}$ where $\hat{\eta}_2(t)$ estimates $\mathrm{E}[g\{Z_{2i}(t)\}] = \eta_2(t)$ and is found by smoothing the data $\big(t, Y_{2i}(t)\big)$ for $i = 1, \ldots, \mathrm{N}$. Next, find the estimator $\hat{S}_{22}(t, t')$ of $S_{22}(t, t') = \mathrm{E}\{Y_{2i}(t)Y_{2i}(t')\} = \mathrm{E}\big[g\{Z_{2i}(t)\}g\{Z_{2i}(t')\}\big]$ by performing bivariate smoothing of the data $\big((t, t'), Y_{2i}(t)Y_{2i}(t')\big)$ for $i = 1, \ldots, \mathrm{N}$, once again removing the diagonals before smoothing. The estimator of the latent process covariance operator for the second response is given by

$$\widetilde{K}_{22}(t, t') = \{\hat{S}_{22}(t, t') - \hat{\eta}_2(t)\hat{\eta}_2(t')\}/[g^{(1)}\{\hat{\mu}_2(t)\}g^{(1)}\{\hat{\mu}_2(t')\}]. \tag{9}$$

Equation (9) was developed for a univariate response, so in order to estimate the latent cross covariance operator $K_{12}(t, t') = K_{21}(t', t) = \mathrm{Cov}\{Z_{1i}(t), Z_{2i}(t')\}$, we must derive an analogous estimator. This requires the following Taylor expansion, also given by equation (5) in Hall et al. (2008),

$$\begin{aligned} g\{Z_i(t)\} =& g\{\mu(t)\} + \delta X_i(t)g^{(1)}\{\mu(t)\} + \frac{1}{2}\delta^2\{X_i(t)\}^2 g^{(2)}\{\mu(t)\} \\ &+ \frac{1}{6}\delta^3\{X_i(t)\}^3 g^{(3)}\{\mu(t)\} + O_p(\delta^4). \end{aligned} \tag{10}$$

We can expand the covariance of the observed processes $\mathrm{Cov}\big\{Y_{1i}(t), Y_{2i}(t')\big\} = \mathrm{Cov}\big[Z_{1i}(t), g\{Z_{2i}(t')\}\big]$ by substituting (10) for

$g\{Z_{2i}(t')\}$ and $\mu_1(t) + \delta X_{1i}(t)$ for $Z_{1i}(t)$, which gives

$$\mathrm{Cov}\big\{Y_{1i}(t), Y_{2i}(t')\big\} = g^{(1)}\{\mu_2(t')\}\mathrm{Cov}\{\delta X_{1i}(t), \delta X_{2i}(t')\} + O(\delta^4). \quad (11)$$

Note that the term (suppressed from equation (11)) $\delta^3 \frac{1}{2} g^{(2)}\{\mu_2(t')\}\mathrm{Cov}\{X_{1i}(t), X_{2i}^2(t')\} = 0$ due to $\mathrm{Cov}\{X_{1i}(t), X_{2i}^2(t')\} = \mathrm{E}\big\{X_{1i}(t)X_{2i}^2(t')\big\} = \mathrm{E}[X_{2i}^2(t')\mathrm{E}\{X_{1i}(t)|X_{2i}(t')\}] = \sigma_1/\sigma_2\rho\mathrm{E}[X_{2i}^3(t')] = 0$ since $X_{1i}(t)|X_{2i}(t') \sim \mathrm{N}\big(\sigma_1/\sigma_2\rho X_{2i}(t'), (1-\rho^2)\sigma_1^2\big)$. Now, because $\mathrm{Cov}\{Z_{1i}(t), Z_{2i}(t')\} = \mathrm{Cov}\{\delta X_{1i}(t), \delta X_{2i}(t')\}$, we have from (11) that $K_{12}(t,t') = \mathrm{Cov}\{Z_{1i}(t), Z_{2i}(t')\} = \mathrm{Cov}\big\{Y_{1i}(t), Y_{2i}(t')\big\}/g^{(1)}\{\mu_2(t')\} + O(\delta^4)$, which, assuming the effect of $O(\delta^4)$ is negligible, leads to

$$K_{12}(t,t') = \mathrm{Cov}\big\{Y_{1i}(t), Y_{2i}(t')\big\}/g^{(1)}\{\mu_2(t')\}. \quad (12)$$

Estimation of (12) requires a smooth estimate $\hat{\eta}_1(t)$ of $\mathrm{E}[Y_{1i}(t)] = \eta_1(t)$ which is found by smoothing the data $\big(t, Y_{1i}(t)\big)$ for $i = 1, \ldots, \mathrm{N}$. We obtain the estimator $\hat{S}_{12}(t,t')$ of $S_{12}(t,t') = \mathrm{E}\{Y_{1i}(t)Y_{2i}(t')\} = \mathrm{E}\big[Y_{1i}(t)g\{Z_{2i}(t')\}\big]$ by performing bivariate smoothing of the data $\big((t,t'), Y_{1i}(t)Y_{2i}(t')\big)$ for $i = 1, \ldots, \mathrm{N}$, removing the diagonals before smoothing. The resulting smooth estimator of the latent cross covariance is

$$\widetilde{K}_{12}(t,t') = \{\hat{S}_{12}(t,t') - \hat{\eta}_1(t)\hat{\eta}_2(t')\}/g^{(1)}\{\hat{\mu}_2(t')\}, \quad (13)$$

which is the direct analogue to (9).

Appendix C

In this section we present the derivations of the conditional posterior distributions.

Random effects

Let L_i be the number of subunits t observed for subject i and define the latent response vector for subject i as $\boldsymbol{W}_i = [W_1(t_1), ..., W_1(t_{\mathrm{L}_i}), W^2(t_1), ..., W^2(t_{\mathrm{L}_i})]^{\mathrm{T}}$, with corresponding mean vector $\mathrm{E}(\boldsymbol{W}_i) = \boldsymbol{U}_i\boldsymbol{\beta} + \boldsymbol{\Psi}_i\boldsymbol{\alpha}_i$. Assume $\boldsymbol{W}_i|\boldsymbol{\alpha}_i \sim \mathrm{N}_{2\mathrm{L}_i}\big(\boldsymbol{U}_i\boldsymbol{\beta} + \boldsymbol{\Psi}_i\boldsymbol{\alpha}_i, \boldsymbol{D}_i\big)$ where $\boldsymbol{D}_i = \mathrm{diag}(\tau_1^2, 1) \otimes \mathbf{I}_{\mathrm{L}_i}$, or in terms of the precision, $\boldsymbol{P}_i = \boldsymbol{D}_i^{-1} = \mathrm{diag}(\omega_1, 1) \otimes \mathbf{I}_{\mathrm{L}_i}$.

Also assume the m × 1 vector $\boldsymbol{\alpha}_i|\boldsymbol{Q} \sim \mathrm{N_m}(0, \boldsymbol{Q}^{-1})$ for $i = 1, \ldots, \mathrm{N}$ and for the m × m covariance matrix $\boldsymbol{Q}^{-1}$, or equivalently, the precision matrix $\boldsymbol{Q}$. Define $\boldsymbol{R}_i = \boldsymbol{W}_i - \boldsymbol{U}_i\boldsymbol{\beta}$. To find the posterior for $\boldsymbol{\alpha}_i|\cdot$ we know

$$p(\boldsymbol{\alpha}_i|\cdot) \propto p(\boldsymbol{W}_i|\cdot) \times p(\boldsymbol{\alpha}_i|\boldsymbol{Q}) \tag{14}$$

$$\propto \exp\left[-\frac{1}{2}\left\{(\boldsymbol{\Psi}_i\boldsymbol{\alpha}_i - \boldsymbol{R}_i)^{\mathrm{T}}\boldsymbol{P}_i(\boldsymbol{\Psi}_i\boldsymbol{\alpha}_i - \boldsymbol{R}_i) + \boldsymbol{\alpha}_i^{\mathrm{T}}\boldsymbol{Q}\boldsymbol{\alpha}_i\right\}\right], \tag{15}$$

$$\propto \exp\left[-\frac{1}{2}\left\{\boldsymbol{\alpha}_i^{\mathrm{T}}\boldsymbol{\Psi}_i^{\mathrm{T}}\boldsymbol{P}_i\boldsymbol{\Psi}_i\boldsymbol{\alpha}_i - 2\boldsymbol{R}_i^{\mathrm{T}}\boldsymbol{P}_i\boldsymbol{\Psi}_i\boldsymbol{\alpha}_i + \boldsymbol{R}_i^{\mathrm{T}}\boldsymbol{P}_i\boldsymbol{R}_i + \boldsymbol{\alpha}_i^{\mathrm{T}}\boldsymbol{Q}\boldsymbol{\alpha}_i\right\}\right], \tag{16}$$

and ignoring terms not involving $\boldsymbol{\alpha}_i$ or $\boldsymbol{Q}$ results in

$$p(\boldsymbol{\alpha}_i|\cdot) \propto \exp\left(-\frac{1}{2}\left[-2\boldsymbol{R}_i^{\mathrm{T}}\boldsymbol{P}_i\boldsymbol{\Psi}_i\boldsymbol{\alpha}_i + \boldsymbol{\alpha}_i^{\mathrm{T}}\left\{\boldsymbol{\Psi}_i^{\mathrm{T}}\boldsymbol{P}_i\boldsymbol{\Psi}_i + \boldsymbol{Q}\right\}\boldsymbol{\alpha}_i\right]\right). \tag{17}$$

We want to form this term into the kernel of a Gaussian distribution where the exponent is $-1/2(\boldsymbol{\alpha}_i - \boldsymbol{M})^{\mathrm{T}}\boldsymbol{V}^{-1}(\boldsymbol{\alpha}_i - \boldsymbol{M}) = -1/2(\boldsymbol{\alpha}_i^{\mathrm{T}}\boldsymbol{V}^{-1}\boldsymbol{\alpha}_i - 2\boldsymbol{M}^{\mathrm{T}}\boldsymbol{V}^{-1}\boldsymbol{\alpha}_i + \boldsymbol{M}^{\mathrm{T}}\boldsymbol{V}^{-1}\boldsymbol{M})$ for some matrices $\boldsymbol{M}$ and $\boldsymbol{V}$. To complete the square, set $\boldsymbol{V} = \{\boldsymbol{\Psi}_i^{\mathrm{T}}\boldsymbol{P}_i\boldsymbol{\Psi}_i + \boldsymbol{Q}\}^{-1}$ and match the coefficients of $\boldsymbol{\alpha}_i$, giving $\boldsymbol{R}_i^{\mathrm{T}}\boldsymbol{P}_i\boldsymbol{\Psi}_i = \boldsymbol{M}^{\mathrm{T}}\boldsymbol{V}^{-1} \implies \boldsymbol{M} = \boldsymbol{V}\boldsymbol{\Psi}_i^{\mathrm{T}}\boldsymbol{P}_i\boldsymbol{R}_i$. Thus, the full conditional posterior for $\boldsymbol{\alpha}_i$ is given by

$$\boldsymbol{\alpha}_i|\cdot \sim \mathrm{N}(\boldsymbol{\mu}_{\boldsymbol{\alpha}}, \boldsymbol{V}_{\boldsymbol{\alpha}})$$

for $\boldsymbol{\mu}_{\boldsymbol{\alpha}} = \{\boldsymbol{\Psi}_i^{\mathrm{T}}\boldsymbol{P}_i\boldsymbol{\Psi}_i + \boldsymbol{Q}\}^{-1}\boldsymbol{\Psi}_i^{\mathrm{T}}\boldsymbol{P}_i(\boldsymbol{W}_i - \boldsymbol{U}_i\boldsymbol{\beta})$ and $\boldsymbol{V}_{\boldsymbol{\alpha}} = \{\boldsymbol{\Psi}_i^{\mathrm{T}}\boldsymbol{P}_i\boldsymbol{\Psi}_i + \boldsymbol{Q}\}^{-1}$.

Random effects precision matrix

Assume the m × 1 vector $\boldsymbol{\alpha}_i|\boldsymbol{Q} \sim \mathrm{N_m}(0, \boldsymbol{Q}^{-1})$ for $i = 1, \ldots, \mathrm{N}$ and the m × m precision matrix $\boldsymbol{Q} \sim \mathrm{Wishart_m}(\boldsymbol{V}, \nu)$ for which the kernel of the density is given by $|\boldsymbol{Q}|^{(\nu-\mathrm{m}-1)/2}\exp\left\{-\frac{1}{2}\mathrm{tr}(\boldsymbol{V}^{-1}\boldsymbol{Q})\right\}$. Define $\boldsymbol{S} = \sum_{i=1}^{\mathrm{N}}\boldsymbol{\alpha}_i\boldsymbol{\alpha}_i^{\mathrm{T}}$ as the sum of squares matrix of $\boldsymbol{\alpha}_i$. We use $\boldsymbol{S}$ to write $\sum_{i=1}^{\mathrm{N}}\boldsymbol{\alpha}_i^{\mathrm{T}}\boldsymbol{Q}\boldsymbol{\alpha}_i = \mathrm{tr}(\sum_{i=1}^{\mathrm{N}}\boldsymbol{\alpha}_i^{\mathrm{T}}\boldsymbol{Q}\boldsymbol{\alpha}_i) = \mathrm{tr}(\sum_{i=1}^{\mathrm{N}}\boldsymbol{\alpha}_i\boldsymbol{\alpha}_i^{\mathrm{T}}\boldsymbol{Q}) = \mathrm{tr}(\boldsymbol{S}\boldsymbol{Q})$ in the kernel of the multivariate normal density, using the properties $\mathrm{tr}(a) = a$ for scalar a and $\mathrm{tr}(\boldsymbol{A}\boldsymbol{B}) = \mathrm{tr}(\boldsymbol{B}\boldsymbol{A})$. We also use the following properties of the trace to combine like-terms: 1) $|\boldsymbol{A}|^{-1} = |\boldsymbol{A}^{-1}|$ for $\boldsymbol{A}$ invertible and 2) for two square matrices $\boldsymbol{A}$ and $\boldsymbol{B}$ of the same dimension, $\mathrm{tr}(\boldsymbol{A}+\boldsymbol{B}) = \mathrm{tr}(\boldsymbol{A})+\mathrm{tr}(\boldsymbol{B})$. Using

this, we can show the conditional posterior $p(\boldsymbol{Q}|\cdot)$ for $\boldsymbol{Q}$ is proportional to the kernel of a $\text{Wishart}_{\text{m}}\{(\boldsymbol{S}+\boldsymbol{V}^{-1})^{-1}, \text{N}+\nu\}$:

$$\begin{aligned} p(\boldsymbol{Q}|\cdot) &\propto \prod_{i=1}^{\text{N}} p(\boldsymbol{\alpha}_i|\boldsymbol{Q}) \times p(\boldsymbol{Q}) \\ &\propto |\boldsymbol{Q}^{-1}|^{-\text{N}/2} \text{exp}\left\{-\frac{1}{2}\text{tr}(\boldsymbol{S}\boldsymbol{Q})\right\} \times |\boldsymbol{Q}|^{(\nu-\text{m}-1)/2} \text{exp}\left\{-\frac{1}{2}\text{tr}(\boldsymbol{V}^{-1}\boldsymbol{Q})\right\} \\ &\propto |\boldsymbol{Q}|^{(\text{N}+\nu-\text{m}-1)/2} \text{exp}\left[-\frac{1}{2}\text{tr}\{(\boldsymbol{S}+\boldsymbol{V}^{-1})\boldsymbol{Q}^{-1}\}\right]. \end{aligned} \tag{18}$$

Thus, $\boldsymbol{Q}|\cdot \sim \text{Wishart}_{\text{m}}\{(\boldsymbol{S}+\boldsymbol{V}^{-1})^{-1}, \text{N}+\nu\}$.

Fixed effects

$$p(\boldsymbol{\beta}|\cdot) \propto \prod_{i=1}^{\text{N}} p(\boldsymbol{W}_i|\cdot) \times p(\boldsymbol{\beta}) \tag{19}$$

$$\propto \text{exp}\left(-\frac{1}{2}\left[\sum_{i=1}^{\text{N}} \left\{(\boldsymbol{U}_i\boldsymbol{\beta}-\boldsymbol{U}_i)^{\text{T}}\boldsymbol{P}_i(\boldsymbol{U}_i\boldsymbol{\beta}-\boldsymbol{U}_i)\right\} + \boldsymbol{\beta}^{\text{T}}\boldsymbol{C}\boldsymbol{\beta}\right]\right) \tag{20}$$

$$\propto \text{exp}\left[-\frac{1}{2}\sum_{i=1}^{\text{N}} \left\{(\boldsymbol{U}_i\boldsymbol{\beta}-\boldsymbol{U}_i)^{\text{T}}\boldsymbol{P}_i(\boldsymbol{U}_i\boldsymbol{\beta}-\boldsymbol{U}_i)\right\} - \frac{1}{2}\boldsymbol{\beta}^{\text{T}}\boldsymbol{C}\boldsymbol{\beta}\right] \tag{21}$$

$$\propto \text{exp}\left[-\frac{1}{2}\sum_{i=1}^{\text{N}} \left\{\boldsymbol{\beta}^{\text{T}}\boldsymbol{U}_i^{\text{T}}\boldsymbol{P}_i\boldsymbol{U}_i\boldsymbol{\beta} - 2\boldsymbol{U}_i^{\text{T}}\boldsymbol{P}_i\boldsymbol{U}_i\boldsymbol{\beta} + \boldsymbol{U}_i^{\text{T}}\boldsymbol{P}_i\boldsymbol{U}_i\right\} - \frac{1}{2}\boldsymbol{\beta}^{\text{T}}\boldsymbol{C}\boldsymbol{\beta}\right] \tag{22}$$

and ignoring constant terms results in

$$\propto \text{exp}\left[-\frac{1}{2}\left\{\boldsymbol{\beta}^{\text{T}}\boldsymbol{C}\boldsymbol{\beta} + \boldsymbol{\beta}^{\text{T}}\left(\sum_{i=1}^{\text{N}}\boldsymbol{U}_i^{\text{T}}\boldsymbol{P}_i\boldsymbol{U}_i\right)\boldsymbol{\beta} - 2\left(\sum_{i=1}^{\text{N}}\boldsymbol{U}_i^{\text{T}}\boldsymbol{P}_i\boldsymbol{U}_i\right)\boldsymbol{\beta}\right\}\right] \tag{23}$$

$$\propto \text{exp}\left[-\frac{1}{2}\left\{\boldsymbol{\beta}^{\text{T}}\left(\boldsymbol{C} + \sum_{i=1}^{\text{N}}\boldsymbol{U}_i^{\text{T}}\boldsymbol{P}_i\boldsymbol{U}_i\right)\boldsymbol{\beta} - 2\left(\sum_{i=1}^{\text{N}}\boldsymbol{U}_i^{\text{T}}\boldsymbol{P}_i\boldsymbol{U}_i\right)\boldsymbol{\beta}\right\}\right] \tag{24}$$

As we did before, we want to form this term into the kernel of a Gaussian distribution and where the exponent is $-1/2(\boldsymbol{\beta} - \boldsymbol{M})^{\mathrm{T}}\boldsymbol{V}^{-1}(\boldsymbol{\beta} - \boldsymbol{M}) = -1/2(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{V}^{-1}\boldsymbol{\beta} - 2\boldsymbol{M}^{\mathrm{T}}\boldsymbol{V}^{-1}\boldsymbol{\beta} + \boldsymbol{M}^{\mathrm{T}}\boldsymbol{V}^{-1}\boldsymbol{M})$ for some matrices $\boldsymbol{M}$ and $\boldsymbol{V}$. To complete the square, set $\boldsymbol{V} = \left(\boldsymbol{C} + \sum_{i=1}^{\mathrm{N}} \boldsymbol{U}_i^{\mathrm{T}}\boldsymbol{P}_i\boldsymbol{U}_i\right)^{-1}$ and match the coefficients of $\boldsymbol{\beta}$, giving $\sum_{i=1}^{\mathrm{N}} \boldsymbol{U}_i^{\mathrm{T}}\boldsymbol{P}_i\boldsymbol{U}_i = \boldsymbol{M}^{\mathrm{T}}\boldsymbol{V}^{-1} \implies \boldsymbol{M} = \boldsymbol{V}\left(\sum_{i=1}^{\mathrm{N}} \boldsymbol{U}_i^{\mathrm{T}}\boldsymbol{P}_i\boldsymbol{U}_i\right)^{\mathrm{T}}$. Thus, the full conditional posterior for $\boldsymbol{\beta}$ is given by

$$\boldsymbol{\beta}|\cdot \sim \mathrm{N}(\boldsymbol{\mu}_{\boldsymbol{\beta}}, \boldsymbol{V}_{\boldsymbol{\beta}})$$

for $\boldsymbol{\mu}_{\boldsymbol{\beta}} = \boldsymbol{V}_{\boldsymbol{\beta}}\left\{\sum_{i=1}^{\mathrm{N}}(\boldsymbol{W}_i - \boldsymbol{\Psi}_i\boldsymbol{\alpha}_i)^{\mathrm{T}}\boldsymbol{P}_i\boldsymbol{U}_i\right\}^{\mathrm{T}}$ and $\boldsymbol{V}_{\boldsymbol{\beta}} = \left(\sigma_{\boldsymbol{\beta}}^{-2}\mathbf{I}_{\mathrm{r}} + \sum_{i=1}^{\mathrm{N}} \boldsymbol{U}_i^{\mathrm{T}}\boldsymbol{P}_i\boldsymbol{U}_i\right)^{-1}$.

Error Variance (Precision)

Assume that the error precision $\omega_1 \sim \mathrm{Gamma}(g, h)$ where we parameterize the density such that g is the shape parameter and $h = 1/s$ is the inverse of the scale parameter, called the rate parameter. Specifically, if $X \sim \mathrm{Gamma}(g, h)$ then $p(x|g, h) = x^{g-1}e^{-xh}\{h^g/\Gamma(g)\}$. For simplicity of notation, denote the continuous response at t_ℓ for subject i as $Y_{i\ell} = Y_{1i}(t_\ell)$ for $i = 1, \ldots, \mathrm{N}$ and $\ell = 1, \ldots, \mathrm{L}_i$, and define the total number of responses observed as $n = \sum_{i=1}^{\mathrm{N}} \mathrm{L}_i$. Let $Y_{i\ell} \overset{indep}{\sim} \mathrm{N}(\mu_{i\ell}, \text{precision} = \omega_1)$. Then

$$p(\omega_1|\cdot) \propto \prod_{i=1}^{\mathrm{N}}\prod_{\ell=1}^{\mathrm{L}_i} p(Y_{i\ell}|\cdot) \times p(\omega_1) \tag{25}$$

$$\propto \omega_1^{n/2} \exp\left\{-\frac{\omega_1}{2}\sum_{i=1}^{\mathrm{N}}\sum_{\ell=1}^{\mathrm{L}_i}(Y_{i\ell} - \mu_{i\ell})^2\right\} \times \omega_1^{g-1}\exp(-\omega_1 h) \tag{26}$$

$$\propto \omega_1^{(n/2+g)-1} \exp\left[-\omega_1\left\{\frac{1}{2}\sum_{i=1}^{\mathrm{N}}\sum_{\ell=1}^{\mathrm{L}_i}(Y_{i\ell} - \mu_{i\ell})^2 + h\right\}\right]. \tag{27}$$

This is the kernel of a Gamma density, so the posterior for $\omega_1|\cdot \sim \mathrm{Gamma}(g_\omega, h_\omega)$ with shape and rate parameters $g_\omega = n/2 + g$ and $h_\omega = 1/2\sum_{i=1}^{\mathrm{N}}\sum_{\ell=1}^{\mathrm{L}_i}(Y_{i\ell} - \mu_{i\ell})^2 + h$.